

multithreading. Performance of such a processor is highly dependent on the accuracy of its branch area was devoted to branch prediction on EV8. The Alpha EV8 branch pre ...

Keywords: EV8 processor, Branch Prediction

5 The effect of instruction fetch bandwidth on value prediction

Freddy Gabbay, Avi Mendelson

April 1998 **ACM SIGARCH Computer Architecture News , Proceedings of the 25th annual international conference on Computer architecture**, Volume 26 Issue 3

Full text available:

 [pdf\(1.32 MB\)](#)  [Publisher Site](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#)

Value prediction attempts to eliminate true-data dependencies by dynamically predicting the outcome of dependent instructions based on that prediction. In this paper we attempt to understand the limits of value prediction. We show that the instruction-fetch bandwidth and the issue rate have a very significant impact on the performance of value prediction. We also study how recent techniques to improve the instruction-fetch bandwidth and the issue rate can be used to improve the performance of value prediction ...

6 Data and memory optimization techniques for embedded systems

P. R. Panda, F. Catthoor, N. D. Dutt, K. Danckaert, E. Brockmeyer, C. Kulkarni, A. Vandercappelle, P. De Dinechin, and J. E. Wilen

April 2001 **ACM Transactions on Design Automation of Electronic Systems (TODAES)**, Volume 6 Issue 1

Full text available:

 [pdf\(339.91 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#)

We present a survey of the state-of-the-art techniques used in performing data and memory-related optimizations. These optimizations are targeted directly or indirectly at the memory subsystem, and impact one or more of the following: memory access time, memory bandwidth, performance, and power dissipation of the resulting implementation. We first examine architecture transformations. We next cover a broad spectrum of optimization techniques ...

Keywords: DRAM, SRAM, address generation, allocation, architecture exploration, code transformation, data flow analysis, design automation, memory architecture customization, memory power dissipation, register file, size estimation

7 Converting thread-level parallelism to instruction-level parallelism via simultaneous multithreading

Jack L. Lo, Joel S. Emer, Henry M. Levy, Rebecca L. Stamm, Dean M. Tullsen, S. J. Eggers

August 1997 **ACM Transactions on Computer Systems (TOCS)**, Volume 15 Issue 3

Full text available:

 [pdf\(526.39 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#)

To achieve high performance, contemporary computer systems rely on two forms of parallelism: instruction-level parallelism (ILP) and thread-level parallelism (TLP). Wide-issue super-scalar processors exploit ILP by executing multiple instructions in parallel. Multiprocessors (MP) exploit TLP by executing different threads in parallel on different processors. However, MP systems must statically partition processor resources, thus preventing threads from sharing processor resources ...

Keywords: cache interference, instruction-level parallelism, multiprocessors, multithreading, simultaneous multithreading

8 A general framework for prefetch scheduling in linked data structures and its application to pointer-chasing

Seungryul Choi, Nicholas Kohout, Sumit Pamnani, Dongkeun Kim, Donald Yeung

May 2004 **ACM Transactions on Computer Systems (TOCS)**, Volume 22 Issue 2

Full text available:

 [pdf\(2.45 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#)

Pointer-chasing applications tend to traverse composite data structures consisting of multiple independent pointer chains. Traversing a single pointer chain leads to the serialization of memory operations, the traversal of independent pointer chains leads to the parallel execution of memory operations. This article investigates exploiting such *interchain memory parallelism* for the purpose of prefetching. Previous work ...

Keywords: Data prefetching, memory parallelism, pointer-chasing code

9 TRIPS: A polymorphous architecture for exploiting ILP, TLP, and DLP

Karthikeyan Sankaralingam, Ramadas Nagarajan, Haiming Liu, Changkyu Kim, Jaehyuk Huh, Nitya Vaidya, and Robert G. McDonald

March 2004 **ACM Transactions on Architecture and Code Optimization (TACO)**, Volume 1 Issue 1

Full text available:

 [pdf\(832.30 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#)

This paper describes the *polymorphous* TRIPS architecture that can be configured for different grain sizes. The architecture is the first in a class of post-RISC, dataflow-like instruction sets called explicit data-granular. It features hardware mechanisms that enable the processing cores and the on-chip memory system to be configured for instruction, data, or thread-level parallelism. To adapt ...

Keywords: Computer architecture, configurable computing, scalable and high-performance computing

10 The white dwarf: a high-performance application-specific processor

A. Wolfe, M. Breternitz, C. Stephens, A. L. Ting, D. B. Kirk, R. P. Bianchini, J. P. Shen

May 1988 **ACM SIGARCH Computer Architecture News, Proceedings of the 15th Annual International Conference on Computer Architecture**, Volume 16 Issue 2

Full text available:  [pdf\(1.40 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

This paper presents the design and implementation of a high-performance special-purpose processor for element analysis algorithms. The White Dwarf CPU contains two Am29325 32-bit floating-point processors. It employs a wide-instruction word architecture in which the application algorithm is directly implemented. It is compatible and interfaces with a SUN 31160 host. The system ...

11 Distributed operating systems

Andrew S. Tanenbaum, Robbert Van Renesse

December 1985 **ACM Computing Surveys (CSUR)**, Volume 17 Issue 4

Full text available:  [pdf\(5.49 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

Distributed operating systems have many aspects in common with centralized ones, but they also have some unique features. This survey provides an introduction to distributed operating systems, and especially to current university research about them. It describes the basic concepts of a distributed operating system and how it is distinguished from a computer network, various key design issues, and the evolution of current research projects are examined in some detail ...

12 Increasing the instruction fetch rate via multiple branch prediction and a branch address cache

Tse-Yu Yeh, Deborah T. Marr, Yale N. Patt

August 1993 **Proceedings of the 7th international conference on Supercomputing**

Full text available:  [pdf\(1.13 MB\)](#)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#), [review](#)

13 A parallel embedded-processor architecture for ATM reassembly

Richard F. Hobson, P. S. Wong

February 1999 **IEEE/ACM Transactions on Networking (TON)**, Volume 7 Issue 1

Full text available:  [pdf\(331.21 KB\)](#)

Additional Information: [full citation](#), [references](#), [citations](#), [index term](#)

Keywords: ATM, embedded systems, medium access control, segmentation and reassembly

14 Embedded applications: AES and the cryptonite crypto processor

Dino Oliva, Rainer Buchty, Nevin Heintze

October 2003 **Proceedings of the 2003 international conference on Compilers, architecture and synthesis for embedded systems**

Full text available:  [pdf\(346.09 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

CRYPTONITE is a programmable processor tailored to the needs of crypto algorithms. The design is based on a methodology of application analysis in which standard crypto algorithms (AES, DES, MD5, SHA-1, etc) were distilled. This methodology and use AES as a central example. Starting with a functional description of AES, we show how to implement AES efficiently in hardware, and present several novel optimizations (which ...)

Keywords: AES, architecture, cryptography, high-bandwidth, high-speed, processor, round key generation

15 Performance evaluation and improvement of a dynamically microprogrammable computer with a microcode cache

Shinji Tomita, Kiyoshi Shibayama, Toshiaki Kitamura, Hiroshi Hagiwara

November 1980 **Proceedings of the 13th annual workshop on Microprogramming**

A new microprogrammable computer with low-level parallelism was built and has been utilized as research-oriented applications such as real-time processings on static/dynamic images, pictures as virtual machines including high (intermediate) level language machines. The design goal of a research is a high degree of processing power and system flexi ...

16 The Vector-Thread Architecture

March 2004

ACM SIGARCH Computer Architecture News , Proceedings of the 31st annual conference on Computer architecture, Volume 32 Issue 2

Full text available:  [pdf\(317.13 KB\)](#)Additional Information: [full citation](#), [abstract](#)

The vector-thread (VT) architectural paradigm unifies the vector and multithreaded compute mode with a control processor and a vector of virtualprocessors (VPs). The control processor can use vector or all the VPs or each VP can use thread-fetches to direct its own control flow. A seamless intermixing allows a VT architecture to flexibly and compactly encode application ...

17 Response Time Analysis of Multiprocessor Computers for Database Support

Roger K. Shultz, Roy J. Zingg

March 1984

ACM Transactions on Database Systems (TODS), Volume 9 Issue 1

Full text available:  [pdf\(2.27 MB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#)

Comparison of three multiprocessor computer architectures for database support is made possible. These expressions are derived by parameterizing algorithms performed by each machine to execute represent properties of the database and components of the machines. Studies of particular parameters of conventional machine technology, for low selectivity, high duplicate occurrence, ...

18 Retrieval operations and data representations in a context-addressed disc system

Stanley Y. W. Su, George P. Copeland, G. Jack Lipovski

November 1973 **Proceedings of the 1973 meeting on Programming languages and informatics**

Full text available:  [pdf\(1.15 MB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#)

This paper attempts to demonstrate that simple expansion of the processing capabilities of fixed data mappings from high-level retrieval language to machine language and from user oriented data representation (storage structure) which are found necessary in conventional von Neumann machines. The built in the disc read and write heads for each disc track allow inf ...

19 System-level power optimization: techniques and tools

Luca Benini, Giovanni de Micheli

April 2000

ACM Transactions on Design Automation of Electronic Systems (TODAES), Volume 5 Issue 1

Full text available:  [pdf\(385.22 KB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#)

This tutorial surveys design methods for energy-efficient system-level design. We consider electro software layers. We consider the three major constituents of hardware that consume energy, namely units, and we review methods of reducing their energy consumption. We also study models for analysis for energy-efficient software design and compilation. This survey ...

20 Exploiting choice: instruction fetch and issue on an implementable simultaneous multithreaded processor

Dean M. Tullsen, Susan J. Eggers, Joel S. Emer, Henry M. Levy, Jack L. Lo, Rebecca L. Stamm

May 1996

ACM SIGARCH Computer Architecture News , Proceedings of the 23rd annual conference on Computer architecture, Volume 24 Issue 2

Full text available:  [pdf\(1.48 MB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#)

Simultaneous multithreading is a technique that permits multiple independent threads to issue multiple instructions in parallel. We demonstrate the performance potential of simultaneous multithreading, based on a somewhat idealized model. The throughput gains from simultaneous multithreading can be achieved without extensive changes to hardware structures or sizes. We present an architecture for s ...